

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	May 23, 1996	Final October 1992-January 1996	
4. TITLE AND SUBTITLE RealTime Communication for Distributed Processing			5. FUNDING NUMBERS F 49620-92-J-0441DEF
6. AUTHOR(S) Douglas S. Reeves			AFOSR-TR-96 O305
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Computer Science North Carolina State University Box 8206 Raleigh, NC 27695-8206			10. SPONSORING/MONITORING AGENCY REPORT NUMBER 19960624 255
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research Bolling Air Force Base, DC 20332-0001			
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT no restrictions		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The subject of this research project was real-time communication in switched networks (the Internet, and ATM). Distributed applications that have real-time requirements include voice and video transmission, real-time simulation, data acquisition, and distributed command and control. Real-time communication protocols should be practical to implement, easy to use, adaptable, and make efficient use of network resources. Previous research on hard and soft real-time communication was summarized and unified. A method which combines the best features of statistical multiplexing and deterministic delivery was then proposed. This method achieves high utilization and provides strong end to end quality of service (QoS) guarantees. A method of dynamic resource allocation was proposed and evaluated. This method uses less resources, is significantly easier to use, and provides equally good QoS, compared with the best static methods. Routing algorithms for uni- and multi-casting of real-time data were also investigated. The limitations of previous algorithms were shown, and improvements were suggested. A simulator for multicast routing experimentation was developed and made available to other researchers. Statistical modelling of variable bit-rate, compressed video traffic was investigated. Such models are useful for network performance evaluation and management. A method was proposed for shaping and smoothing of VBR video traffic at the user-network interface.			
14. SUBJECT TERMS Networks, Quality of service, Routing, Multicasting, Resource Allocation			15. NUMBER OF PAGES 21
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT none

Real-Time Communication for Distributed Computing
AFOSR Grant F49620-92-J-0441DEF

FINAL TECHNICAL REPORT

October 1992 – January 1996

Douglas S. Reeves

Department of Computer Science
N. C. State University
Box 8206
Raleigh, NC 27695-8206
Telephone: (919) 515-2044
Fax: (919) 515-7896

Email: reeves@eos.ncsu.edu

1 ABSTRACT

The subject of this research project was real-time communication in switched networks (the Internet, and ATM). Distributed applications that have real-time requirements include voice and video transmission, real-time simulation, data acquisition, and distributed command and control. Real-time communication protocols should be practical to implement, easy to use, adaptable, and make efficient use of network resources.

Previous research on hard and soft real-time communication was summarized and unified. A method which combines the best features of statistical multiplexing and deterministic delivery was then proposed. This method achieves high utilization and provides strong end-to-end quality of service (QoS) guarantees. A method of dynamic resource allocation was proposed and evaluated. This method uses less resources, is significantly easier to use, and provides equally good QoS, compared with the best static methods.

Routing algorithms for uni- and multi-casting of real-time data were also investigated. The limitations of previous algorithms were shown, and improvements were suggested. A simulator for multicast routing experimentation was developed and made available to other researchers.

Statistical modelling of variable bit-rate, compressed video traffic was investigated. Such models are useful for network performance evaluation and management. A method was proposed for shaping and smoothing of VBR video traffic at the user-network interface.

2 RESEARCH OBJECTIVES

Our research is on real-time communication in point-to-point networks (i.e. wide-area networks, switched local-area networks, and multiprocessor interconnection networks). Communication must be real-time for distributed applications that have strict timing requirements. Real-time applications require bounds on the maximum end-to-end delay, delay variation, and packet loss rates in the network. Distributed real-time applications of interest to the Air Force include: voice and video transmission; distributed real-time simulation; data acquisition from distributed sensors; remote command and control; and many others. A substantial percentage of traffic on most networks is expected to be real-time in nature. For cost reasons, it is desirable that real-time service be as efficient as possible, and be available on the most widespread networking technologies, like ATM and the Internet. We have proposed network policies and mechanisms which provide strong real-time guarantees. These policies are practical to implement, easy to use, adaptable to many types of applications, and make efficient use of network resources.

Our specific objectives for this research were the following:

- G1 Unify and summarize previous work on real-time communication in packet-switched networks, and promote this as a worthwhile subject in the research community.
- G2 Develop admission control, bandwidth allocation, and packet queueing policies that have the desired properties (practical, usable, adaptable, efficient). Our overall goal is to achieve strong Quality-of-Service (QoS) guarantees, with high network utilization.
- G3 Investigate the suitability of known unicast routing algorithms for real-time traffic, and the interaction between admission control and routing. Develop new methods where needed.

- G4 Investigate the suitability of known multicast routing algorithms for real-time traffic. Determine whether the quality of these methods is good, and whether they are practical to use in large networks. Devise algorithms for near-optimal selection of *cores*, which are special nodes that distribute data to all the members of a multicast group.
- G5 Investigate the transmission of compressed video, which is one of the major applications for real-time communication. Determine whether statistical models of compressed video are practical and accurate, and show how to use to transmit compressed video efficiently, and with acceptable quality.

Our target technologies for this research are the Internet and ATM.

3 STATUS AND ACCOMPLISHMENTS

- G1 Conducted an in-depth survey of work on real-time communication, in collaboration with researchers from the University of Massachusetts, and published this survey. Identified many important open problems and promising directions for future work. Organized and conducted a panel session at the major annual real-time systems symposium to raise interest in this topic. Guest edited a special issue of an IEEE magazine on parallel and distributed real-time computing. Publications [J2,E1].
- G2 Improved an existing method of hard real-time communication proposed by Ferrari. The improvements were due to the use of preemption, cut-through, and adaptive slack allocation. Report [T6] and thesis [T7].
- We proposed a method for real-time communication which provides strong end-to-end guarantees, yet achieves high network utilization. We proposed that flows (of real-time data) be statistically multiplexed at the entrance to the network, but deterministically transmitted inside the network. The result is predictable end-to-end losses and delays (which statistical methods alone will not provide), along with high utilizations (which deterministic methods alone will not provide). We showed how to allocate buffer space optimally for such a scheme. Publication [C2] and thesis [T3].
- We then invented a new method for allocating resources (bandwidth, buffer space, queueing priority) dynamically to real-time applications. This method is based on measuring the actual QoS, and adjusting resources dynamically as required to meet user requirements. For traffic sources whose statistical mean is stationary, the method uses the minimum resources possible to achieve the desired QoS. The method does not require users to provide a statistical model of the application traffic, and is applicable to both ATM and the Internet. We are now in the process of demonstrating its effectiveness on actual traffic traces. Publication [C8], report [T2], and thesis [T3].
- G3 We conducted an in-depth investigation of unicast routing algorithms for real-time data. This work illustrated many short-comings in algorithms which had previously been proposed. We suggested a new method which performs from 5% to 50% better than the

best known techniques. This method considers the constraints of the admission control algorithm while finding a route. Publications [J1,C5] and thesis [T3].

We later proposed another unicasting algorithm which has excellent performance (lowest cost, and meeting delay constraints), and which is very compatible with existing Internet routing protocols. Publication [C6].

G4 We investigated multicast routing algorithms for real-time data in high-speed networks. We showed that delay-constrained heuristics are close to optimal (although expensive to compute), while unconstrained methods are not very good (and a method being promoted by the Internet Engineering Task Force is particularly bad). We devised a routing algorithm for constructing delay-constrained minimum spanning trees. The multicast routing simulator developed in our work has been distributed to over a dozen sites for use in other research. Publications [J3, C1, C7] and report [T4].

Another objective of our research is to identify an optimal set of "cores" or "rendezvous points". The networking community agrees that multicasting in large networks requires the use of cores, yet there has been little work on how to select the best set of cores. This work will have application to ATM and the Internet, and will be useful for both real-time and non-real-time traffic. Work-in-progress (not yet submitted), and report [T1].

G5 We investigated statistical models for variable-bit-rate MPEG-compressed video, using actual video traces. Our results indicate serious shortcomings of some of the proposed methods, and show that a slice-based model is significantly better than a frame-based model. We also conclude that further improvements in statistical modelling must consider scene changes and long-range dependence. We implemented a software tool for parsing and analyzing videos which has been made available to a research group at IBM, and to other researchers. Publications [J4, C3, C4].

We are now investigating methods of reliable and efficient video transmission. Our specific focus is on traffic shaping and policing of compressed video at the entrance to the network. Work-in-progress (not yet submitted).

4 PUBLICATIONS

All of the publications cited below are available both in abstract form and as postscript papers from URL

<ftp://ftp.csc.ncsu.edu/pub/rtcomm/rtcomm.html>

The abstracts are also included at the end of this report.

Edited

E1 D. Reeves and K. Shin, editors, "Special Issue on Parallel and Distributed Real-Time Computing", *IEEE Parallel and Distributed Technology*, Winter 1994.

Published in Journals

- J1 S. Rampal, and D. Reeves. "Routing and Admission Control Algorithms for Multimedia Traffic", *Computer Communications*, North-Holland Publ. Co., Vol. 18, No. 10, October 1995, pp. 755-768.
- J2 C. Aras, J. Kurose, D. Reeves, H. Schulzrinne, "Real-Time Communication in Packet-Switched Networks", *Proceedings of the IEEE*, Vol. 82, Number 1 (January 1994), pp. 122 -139.

Submitted to Journals, in Review

- J3 H. Salama, D. Reeves, and Y. Viniotis, "Evaluation of Multicast Routing Algorithms for Real-Time Communication on High-Speed Networks", submitted to *Journal on Selected Areas of Communication*, January 1996.
- J4 M. Izquierdo and D. Reeves. "Statistical Models of Variable-Bit-Rate Compressed Video", submitted to *Multimedia Systems Journal*, May 1996.

Published in Conferences

- C1 H. Salama, D. Reeves, Y. Viniotis, and T. Sheu. "Evaluation of Multicast Routing Algorithms For Distributed Real-Time Applications in High-Speed Networks", *6th IFIP Conf. on High-Speed Networks*, Chapman-Hall Publ., September 1995.
- C2 S. Rampal, D. Reeves, and D. Agrawal. "End-to-End Guaranteed QoS with Statistical Multiplexing for ATM Networks", in *Performance Modelling and Evaluation of ATM Networks*, ed. D. Kouvatsos, Chapman and Hall Publ., 1995.
- C3 M. Izquierdo and D. Reeves. "Issues Related to Slice-Based Modeling of MPEG VBR-Encoded Video", in *Proc. of IEEE Southeastcon*, Raleigh NC, March 1995.
- C4 M. Izquierdo and D. Reeves. "Statistical Characterization of MPEG VBR-Encoded Video at the Slice Layer", in *Proc. of the SPIE Conf. on Multimedia Computing and Networking*, SanJose CA, March 1995.
- C5 S. Rampal, D. Reeves, and D. P. Agrawal. "An Evaluation of Routing and Admission Control Algorithms for Real-Time Traffic in Packet-Switched Networks", in *Proc. of the 5th IFIP Conf. on High-Performance Networking (HPN '94)*, North-Holland Publ. Co, 1995.

Submitted To Conferences, in Review

- C6 H. Salama, D. Reeves, and Y. Viniotis, "A Distributed Algorithm for Delay-Constrained Unicast Routing", submitted to IEEE Infocom '97, May 1996.
- C7 H. Salama, D. Reeves, and Y. Viniotis, "An Efficient Delay-Constrained Minimum Spanning Tree Heuristic", submitted to Intl. Conf. on Computer Communications and Networks, March 1996.

- C8 S. Rampal, D. Reeves, and Y. Viniotis, "Dynamic Resource Allocation Based on Measured Quality of Service", submitted to Intl. Conf. on Computer Communications and Networks, March 1996.

Technical Reports and Theses

- T1 H. Salama, "The Core Selection Problem for Multicasting of Real-Time Data", Center for Advanced Computing and Communications, N. C. State University, Technical Report, May 1996.
- T2 S. Rampal, D. Reeves, and Y. Viniotis, "Dynamic Resource Allocation for Quality of Service". Center for Advanced Computing and Communications, N. C. State University, TR 96-2, January 1996.
- T3 S. Rampal, "Routing and End-to-End Quality of Service in Multimedia Networks", PhD Thesis, Department of Electrical and Computer Engineering, N. C. State University, August 1995.
- T4 H. Salama, Y. Viniotis, D. Reeves, and T. Sheu, "Multicast Routing Algorithms for High-Speed Networks", IBM Technical Report, Sept. 1994.
- T5 S. Rampal, D. Reeves, and D. Agrawal, "Processor Scheduling Algorithms for Minimizing Buffer Requirements in Multimedia Applications", Center for Communications and Signal Processing, N. C. State University, TR 94-16, July 1994.
- T6 C. Aras, D. Reeves, and R. Luo, "Low Latency, High Acceptance Real-Time Communication in Wide Area Networks", CCSP Technical Report, N. C. State University, March 1993.
- T7 C. Aras, *Communication Networks for Autonomous Mobile Robot Computer Architectures and Distributed Real-Time Computing*, PhD Thesis, Department of Electrical and Computer Engineering, N. C. State University, November 1992.

5 PERSONNEL

The following personnel have been associated with this research.

- Douglas S. Reeves, Principal Investigator. Associate Professor of Computer Science, N. C. State University.
- Dharma Agrawal (Professor) and Yannis Viniotis (Associate Professor), N. C. State University.
- Jim Kurose (Associate Professor), University of Massachusetts.
- Henning Schulzrinne (member of technical staff), GMD-FOKUS (Germany).
- Caglan Aras, now with IBM Corp. at Research Triangle Park, NC. Received his PhD in Computer Engineering from N. C. State in November 1992. Title of thesis: "Communication Networks for Autonomous Mobile Robot Computer Architectures and Distributed Real-Time Computing".

- Sanjeev Rampal, now with IBM Corp. at Research Triangle Park, NC. Received his PhD in Computer Engineering from N. C. State in August 1995. Thesis title: "Routing and End-to-End Quality of Service in Multimedia Networks".
- Hussein Salama, currently pursuing his PhD in Computer Engineering from N. C. State. Thesis topic: "Multicast Routing and Core Selection for Real-Time Traffic". Expected completion date is December 1996.
- Mike Izquierdo, currently pursuing his PhD in Computer Engineering from N. C. State. Thesis topic: "Statistical Modelling and Effective Transmission of Compressed Video". Expected completion date is December 1996.
- Errin Fulp, currently pursuing his PhD in Computer Engineering from N. C. State. Thesis topic: "Dynamic Resource Allocation for Quality of Service". Just beginning his PhD research.
- Two other Ph.D. students, Steven Wright and Herbert Rivera-Sanchez, have just started looking at research problems related to real-time communication. They have not yet picked a topic.

6 PRESENTATIONS

Presentations by D. Reeves

1. "Dynamic Resource Allocation for Quality of Service", annual meeting of the Center for Advanced Computing and Communications (Raleigh, NC), October 1995.
2. "Multicasting for Real-Time Traffic", meeting between AFOSR and DISA personnel (Reston, VA), October 1995.
3. "Evaluation of Multicasting Routing Algorithms for Real-Time Applications", IFIP 6th Conference on High Performance Networking (Mallorca, Spain), September 1995.
4. "Dynamic Resource Allocation for Quality of Service", and "Introduction to Multimedia Technology", AF Rome Labs (Rome, NY), June 1995.
5. "Research in Real-Time Communication", AFOSR Contractor's meeting (Raleigh, NC), March 1995.
6. "End-to-End Guaranteed QoS with Statistical Multiplexing for ATM Networks", AFOSR Contractor's meeting (Washington, DC), September 1994.
7. "End-to-End Guaranteed QoS with Statistical Multiplexing for ATM Networks", IFIP Workshop on Modelling and Evaluation of ATM Networks (Bradford, England), July 1994.
8. "An Evaluation of Routing and Admission Control Algorithms for Real-Time Traffic in Packet-Switched Networks", 5th IFIP Conf. on High-Performance Networking (Grenoble, France), June 1994.

Presentations by Students

1. M. Izquierdo, "Statistical Modelling of MPEG-Compressed Video", IBM Corp. (Research Triangle, NC), April 1995.
2. M. Izquierdo, "Statistical Characterization of MPEG VBR-Encoded Video at the Slice Layer", SPIE Conf. on Multimedia Computing and Networking (SanJose CA), March 1995.
3. M. Izquierdo, "Issues Related to Slice-Based Modeling of MPEG VBR-Encoded Video", IEEE Southeastcon (Raleigh, NC), March 1995.
4. H. Salama, "Multicast Routing for Real-Time Communication", IBM Corp. (Research Triangle, NC), September 1994.

7 CONSULTING, ADVISING, AND TRANSITIONS

Professional Activities

- Organized panel session on "Real-Time Communication" at 12th Real-Time Systems Symposium, December 1993.
- Reviewer of several proposals for both NSF and AFOSR.
- Reviewer for 6 conferences, and more than 10 journals.
- Member of program committee for 3 conferences.
- Associate Editor-in-Chief and Editorial Board Member for *IEEE Parallel and Distributed Technology*.

Interaction with Air Force Personnel

- Participated in joint DISA / AFOSR meeting, October 1995.
- Visited Rome Labs and gave talks, June 1995.
- Participated in 2 AFOSR Contractor's meetings (October 1994 and March 1995).
- Participated in meeting with Rome Labs personnel, November 1994.

Source Code Distributed

- Distributed multicast routing simulator to approximately a dozen users around the world. Available by anonymous ftp access from [ftp.csc.ncsu.edu](ftp://ftp.csc.ncsu.edu), directory pub/rtcomm.
- Provided tool for parsing and displaying relevant information about MPEG-2 compressed videos, at the Transport Stream level (i.e., as accepted by the network). Used at IBM, and provided to approximately a half-dozen other requestors. Also available from the same ftp site.

8 PATENTS

None filed at this time.

9 HONORS AND AWARDS

Currently nominated for Outstanding Research Award, College of Engineering, N. C. State University.

10 CONCLUSION

The Air Force uses an abundance of networks at present, each with its own advantages. These include local area networks, the Internet, microwave and satellite links, leased lines, the phone system, short-range wireless, narrowband ISDN, and others. This abundance of technologies is expensive, difficult to maintain and interoperate, and creates major barriers to developing applications that work throughout the entire enterprise. As a result, the Air Force will benefit a great deal from the development of integrated networks that provide a range of capabilities. Integrated networks must support real-time data. The two major solutions for integrated networking are: extension of the Internet protocols to support real-time traffic; and, Asynchronous Transfer Mode (ATM) technology.

The "Holy Grail" of research in real-time communication is how to provide strong end-to-end guarantees and still achieve high network utilization. These end-to-end guarantees (on maximum packet delay and maximum loss rates) are easy to achieve if you take a very conservative approach to network design. However, 30% network utilization seems unacceptably low. In addition, much of the work on real-time communication assumes that user traffic is static and can be very precisely described. Both of these assumptions seem doubtful. I am ambitious: I believe the Holy Grail can be attained, and I reject these limiting assumptions. Our work in (uni- and multi-casting) routing, admission control, resource allocation, and traffic shaping has helped us measure the distance to that goal more clearly than ever, and has substantially reduced that distance in my opinion.

I am fortunate to be working at the interface between the disciplines of networking and real-time systems. It's an exciting place to be, and few people are willing or able to cross over between these two fields. The continued support of the Air Force has made our work possible, and I'm deeply grateful. I hope the results will justify that support.

11 COPIES OF PUBLICATIONS

All of the publications resulting from this research, including theses and technical reports, are available in both abstract form and as postscript papers from the following URL:

<ftp://ftp.csc.ncsu.edu/pub/rtcomm/rtcomm.html>

The online list of articles, abstracts, and source code is reproduced in the following pages.

Real-Time Communication Project

Source Code

- The Multicast Routing Simulator
- The MPEG-2 Transport Stream Browser

Papers on Resource Allocation and Admission Control

- "Dynamic Resource Allocation for Quality of Service", by Rampal, Reeves, and Viniotis. Center for Advanced Computing and Communications, N. C. State University, TR 96-2, January 1996.
 - Abstract
 - Postscript paper
- "Routing and End-to-End Quality of Service in Multimedia Networks", by S. Rampal. *PhD Thesis*, N. C. State University, August 1995.
 - Abstract
 - Postscript paper
- "Guaranteed End-to-End QOS with Statistical Methods in ATM Networks", by Rampal, Reeves, and Agrawal. *Performance Modelling and Evaluation of ATM Networks*, D. Kouvatsos ed., Chapman and Hall Publ., 1995.
 - Abstract
 - Postscript paper
- "An Approach Towards End-to-End QoS with Statistical Multiplexing in ATM Networks", by Rampal, Reeves, Viniotis, and Agrawal. Technical Report 95/2, Center for Communications and Signal Processing, N. C. State University, January 1995.
 - Abstract
 - Postscript paper
- "Real-Time Communication in Packet-Switched Networks", Aras, Kurose, Reeves, Schulzrinne. *Proc. of the IEEE*, Vol. 82 No. 1, Jan. 1994, pp. 122-139.
 - Abstract
 - Postscript paper
- "Low Latency, High Acceptance Real-Time Communication in Wide Area Networks" by Aras, Reeves, and Luo. Unpublished paper.
 - Abstract
 - Postscript paper

Papers on Multicasting and Routing

- "An Efficient Delay-Constrained Minimum Spanning Tree Heuristic", by Salama, Reeves, and Viniotis, *submitted for publication*, March 1996.
 - Abstract
 - Postscript paper
- "Evaluation of Multicast Routing Algorithms for Real-Time Communication on High-Speed Networks", by Salama, Viniotis, Reeves, *submitted for publication*, January 1996.
 - Abstract

- Postscript paper
- "Routing Algorithms for Multimedia Data", by Rampal and Reeves. *Computer Communications*, North-Holland Publ., October 1995.
 - Abstract
 - Postscript paper
- "Evaluation of Multicast Routing Algorithms for Distributed Real-Time Applications of High-Speed Networks", by Salama, Viniotis, Reeves, and Sheu. *Proc. of 6th IFIP Conf. on High-Performance Networks (HPN '95)*, Sept. 1995.
 - Abstract
 - Postscript paper
- "Multicast Routing Algorithms for High-Speed Networks", by Salama, Viniotis, Reeves, and Sheu. IBM Technical Report, Sept. 1994.
 - Abstract
 - Postscript paper
- "An Evaluation of Routing and Admission Control Algorithms for Real-Time Traffic in Packet-Switched Networks", by Rampal, Reeves, and Agrawal. *Proc. of the 5th IFIP Conf. on High Performance Networks (HPN '94)*, Chapman and Hall, 1995.
 - Abstract
 - Postscript paper
- "Comparison of Multicast Routing Algorithms for High-Speed Networks", by Salama, Viniotis, Reeves, and Sheu. IBM Technical Report.
 - Abstract
 - Compressed Postscript paper

Papers on Video Modelling

- "Issues Related to Slice-Based Modelling of MPEG VBR-Encoded Video", by Izquierdo and Reeves. *Proc. of IEEE Southeastcon '95*, March 1995, Raleigh NC.
 - Abstract
 - Postscript paper
- "Statistical Characterization of MPEG VBR Video at the Slice Layer", by Izquierdo and Reeves. *Proc. of the Conf. on Multimedia Computing and Networking*, SPIE, Feb. 1995, San Jose CA.
 - Abstract
 - Postscript paper

Other Papers on Real-Time Communication

- "Processor Scheduling Algorithms for Minimizing Buffer Requirements in Multimedia Applications". Rampal, Reeves, and Agrawal. Technical Report TR94-16, Center for Communications and Signal Processing, N. C. State University, July 1994.
 - Abstract
 - Postscript paper
 - *Communication Networks for Autonomous Mobile Robot Computer Architectures and Distributed Real-Time Computing*. PhD thesis of Caglan Aras, November 1992, N. C. State University.
 - Abstract
 - Postscript paper
-

Dr. D. S. Reeves, reeves@eos.ncsu.edu

Real-Time Communication Project: Abstracts of Papers

An Efficient Delay-Constrained Minimum Spanning Tree Heuristic

We formulate the problem of constructing broadcast trees for real-time traffic with delay constraints in networks with asymmetric link loads as a delay-constrained minimum spanning tree (DCMST) problem in directed networks. Then we prove that this problem is $\{\text{NP}\}$ -complete, and we propose an efficient heuristic to solve the problem based on Prim's algorithm for the unconstrained minimum spanning tree problem. This is the first heuristic designed specifically for solving the DCMST problem. Simulation results under realistic networking conditions show that our heuristic's performance is close to optimal when the link loads are symmetric as well as when asymmetric link loads are used. Delay-constrained minimum Steiner tree heuristics can be used to solve the DCMST problem. Simulation results indicate that the fastest delay-constrained minimum Steiner tree heuristic, DMCT, is not as efficient as the heuristic we propose, while the most efficient delay-constrained minimum Steiner tree heuristic, BSMA, is much slower than our proposed heuristic and does not construct cheaper delay-constrained broadcast trees.

"Dynamic Resource Allocation for Quality of Service"

Most methods for guaranteeing quality of service (QoS) in packet-switched networks require a static characterization of the user's traffic, and allocate resources accordingly. This process is inconvenient for unsophisticated users, and can result in overallocation of resources. We address the problem of automatically determining the minimal resources necessary to satisfy a specified Quality of Service (QoS) measure using dynamic techniques. Specifically, we study a dynamic control algorithm which determines the minimum bandwidth needed to satisfy a specified average cell loss probability for a given source. This algorithm is called REQS, which stands for Resource Efficient Quality of Service. REQS measures the actual cell losses, and uses this information to dynamically vary the bandwidth allocation. Experimental results show the algorithm converges quickly to the desired solution. The effect of the measurement frequency, source burstiness, buffer size and loss specification on the convergence time of the algorithm are all investigated. The algorithm is robust, converging quickly and accurately under most conditions. We demonstrate the bandwidth savings attainable by this approach over other techniques, such as the equivalent capacity. The applicability of the method for other QoS measures (such as queuing delay percentiles) is also investigated. Some of the applications of this technique include control of rate-enforcing servers, traffic shapers, and call admission control in ATM networks.

" Real-Time Communication in Packet-Switched Networks"

The dramatically increased bandwidths and processing capabilities of future high-speed networks make

possible many distributed real-time applications, such as sensor-based applications and multimedia services. Since these applications will have traffic characteristics and performance requirements that differ dramatically from those of current data-oriented applications, new communication network architectures and protocols will be required. In this paper we discuss the performance requirements and traffic characteristics of various real-time applications, survey recent developments in the areas of network architecture and protocols for supporting real-time services, and develop frameworks in which these, and future, research efforts can be considered.

"An Evaluation of Routing and Admission Control Algorithms for Real-Time Traffic in Packet-Switched Networks"

Networks supporting real-time traffic employ controlled admission of calls to ensure an acceptable quality of service. In this work, we investigate how routing and admission control interact in determining overall performance. We propose and evaluate the use of routing algorithms which take into account the constraints imposed by the admission control algorithms. These algorithms are of the least loaded path type and are found to perform better than sequential routing type algorithms which have been suggested elsewhere. We show that the new algorithms decrease call blocking probability and increase network utilization. The amount of improvement depends on such factors as the admission control function, the traffic mix, and the QOS constraints.

Two deterministic methods of Call Admission Control: Earliest Due Date and Stop&Go are evaluated. The effect of lossy source traffic shaping on routing is studied. The interaction and relative importance of routing, admission control, and traffic shaping is examined.

"Routing Algorithms for Multimedia Data"

Interactive voice and video applications (e.g. teleconferencing) over multi-hop packet-switched networks require bounds on delay, loss, and jitter. This paper compares routing algorithms and admission control methods which are intended to provide the quality of service required by such applications. The method of comparison is detailed simulation. As part of this work, we propose several novel real-time routing algorithms.

The best-performing routing algorithm is a new method which takes into account the constraints imposed by the admission control algorithm. A routing algorithm similar to those used in the phone system is the worst of the dynamic algorithms; static routing is far worse than any of these. The differences between the routing algorithms are greatest when the network is highly connected, and when admission control constraints are difficult to meet. The issue of {lem fairness} is also considered in the evaluation of the routing algorithms.

Three admission control algorithms were compared, under varying traffic conditions and QoS requirements. The impact of traffic shaping on link utilizations was investigated. Overall, the choice of admission control method was more important than the choice of routing algorithm.

This work is the first detailed evaluation of routing algorithms for voice and video traffic in multi-hop

networks that provide strong end-to-end guarantees.

"Comparison of Multicast Routing Algorithms for High-Speed Networks"

Multicast routing algorithms capable of satisfying the quality of service (QoS) requirements of real-time applications will be essential for future high-speed networks. We compare the performance of all of the important MC routing algorithms when applied to networks with asymmetric link loads. Each algorithm is judged based on the quality of the MC trees it generates and its efficiency in managing the network resources. Simulation results over random networks show that unconstrained algorithms are not capable of fulfilling the QoS requirements of real-time applications in wide-area networks. The simulations also reveal that one of the unconstrained algorithms, reverse path multicasting (RPM), is quite inefficient when applied to asymmetric networks. We study how combining routing with resource reservation and admission control in a single module considerably improves RPM's efficiency in managing the network resources. Four delay-constrained algorithms are also studied: three constrained Steiner tree (CST) heuristics and one constrained shortest path tree (CSPT) heuristic. Simulations show that all three CST heuristics behave similarly and that they can manage the network efficiently and construct low cost MC trees that satisfy the QoS requirements of real-time traffic. The execution times of the CST heuristics are considerably larger than those of the unconstrained algorithms. The CSPT heuristic is not as efficient as the CST heuristics, but its execution times are on the same order as the execution times of the unconstrained algorithms.

"Multicast Routing Algorithms for High-Speed Networks"

Distributed Real-time applications requiring quality of service (QoS) guarantees and involving groups of users will dominate future's high speed networks. Efficient multicast routing (MCR) algorithms are necessary for these applications. We survey the existing MCR algorithms and study their characteristics. These algorithms were not designed to provide QoS guarantees. We formulate the multicasting problem for applications with QoS requirements and define the criteria for evaluating the performance of the MCR algorithms when applied to this problem. We also describe the simulator we used for evaluating these algorithms.

"Comparison of Multicast Routing Algorithms for High-Speed Networks"

Distributed multimedia applications requiring quality of service (QoS) guarantees and involving groups of users will dominate future's high-speed networks. Efficient multicast routing algorithms are necessary for these applications. This report compares the performance of multicast routing algorithms when applied to real-time networks. We briefly discuss previous work on multicast routing, then we formulate the multicasting problem for applications with QoS requirements and define the criteria for evaluating the performance of the MC routing algorithms. We study five MC routing algorithms, three unconstrained algorithms, and two heuristics designed specifically for ATM networks. A simulator is used to evaluate the performance of the algorithms. We present simulation results over random network graphs that show that unconstrained routing algorithms are capable of satisfying the delay requirements

of real-time services when applied to networks spanning distances of up to a 1000 Km. We also use simulation to evaluate the ability of the algorithm to balance the network load.

"Evaluation of Multicast Routing Algorithms for Distributed Real-Time Applications on High-Speed Networks"

Multicast (MC) routing algorithms capable of satisfying the QoS requirements of multimedia applications will be essential for future's high-speed networks. We compare the performance of selected MC routing algorithms when applied to networks with asymmetric link loads. Each algorithm is judged based on the quality of the MC tree it generates and its efficiency in managing the network resources. Simulation results over random networks show that unconstrained algorithms are not capable of fulfilling the QoS requirements of applications utilizing networks spanning large areas. One algorithm, reverse path multicasting, is not suitable for asymmetric networks irrespective of the requirements of the application. Three constrained Steiner tree (CST) heuristics are also studied. Simulations show that all three behave similarly and that they can manage the network efficiently and construct low cost MC trees that satisfy the QoS requirements of multimedia traffic. The execution times of the CST heuristics are larger than those of unconstrained algorithms.

"Guaranteed End-to-End QOS with Statistical Methods in ATM Networks"

We investigate a method for supporting diverse quality-of-service requirements in broadband networks based on ATM technology. The method uses deterministic bandwidth reservation at the Virtual Path (VP) level and statistical multiplexing within each VP. A deterministic server such as a Weighted Round Robin (WRR) server is used to enforce bandwidth reservations among the VPs. We develop a connection admission algorithm which accounts for end-to-end delay and loss guarantees for Virtual Circuits which traverse a single VP. We show that under certain conditions the amount of network bandwidth required by a VP is minimized by incurring all the allowable loss at the first link of a VP. Achievable utilization is demonstrated using simulation. The effect of the parameters of the WRR server ({\em i.e.}, the vacation time) on the cell loss probability is also studied using simulation.

"An Approach Towards End-to-End QoS with Statistical Multiplexing in ATM Networks"

We address the problem of providing quality-of-service (QoS) guarantees in a multiple hop packet/cell switched environment while providing high link utilization in the presence of bursty traffic. A scheme based on bandwidth and buffer reservations at the Virtual Path level is proposed for ATM networks. This approach enables us to provide accurate end-to-end QoS guarantees while achieving high utilization by employing statistical multiplexing and traffic shaping of bursty traffic sources. A simple round robin scheduler is proposed for realizing this approach and is shown to be implementable using standard ATM hardware viz. cell spacers. The problem of distributing the bandwidth and buffer space assigned to a VP over its multiple hops is addressed. We prove the optimality of the approach of

allowing all the end-to-end loss to occur at the first hop under some conditions and show that its performance can be bounded with respect to the optimal in other conditions. This results in an equal amount of bandwidth to a VP at each hop and essentially no queueing after the first hop. Using simulations, the average case performance of this approach is also found to be good. Additional simulation results are presented to evaluate the proposed approach.

"Issues Related to Slice-Based Modelling of MPEG VBR-Encoded Video"

This paper presents a slice based model for VBR-encoded MPEG videos. This model's random variable is independent and would fit classical distributions. We analyzed four MPEG-1 VBR encoded high quality, color video sequences. None of the sequences contained audio. We present three types of slice based models and discuss the merits of each. We show the distributions given by each of the models and show their fit to the Gamma and Pareto distributions using the QQ plot.

"Statistical Characterization of MPEG VBR Video at the Slice Layer"

In this paper, we statistically characterized four VBR encoded video sequences, containing I/B/P frames, at the Slice layer with the goal of developing an accurate source model to better understand the bit-rate behavior of these sources. We presented the cells/slice distributions and showed that it is "heavy tailed" and fits the Pareto distribution better than Gamma. We showed that an 8-state Markov Chain fits the cells/slice distribution well, reaching steady state after 37 to 80 transitions (2 to 5 frames). We also showed that the autocorrelation function is quasi-periodic which is mostly due to the frame sequence pattern rather than spatial content. We discussed the impact of I/B/P sequences on multiplexing and dynamic bandwidth allocation and proposed a multiplexing method called Time Shifted Multiplexing (TSM); whereby, the multiplexer attempts to overlap I and P frames of one video stream with B frames of another. This tends to reduce both Peak-to-Mean-Ratio and Coefficient-of-Variation of the multiplexed output stream. We showed that coefficient-of-variation reduced in half and bandwidth requirements reduced by 41% using TSM.

"Processor Scheduling Algorithms for Minimizing Buffer Requirements in Multimedia Applications"

The increasing use of audio, video and other multimedia applications on computers mandates the use of real-time processor scheduling techniques. Real-time scheduling algorithms have mostly concentrated on meeting application deadlines (such as those required for video playback). An important requirement of such applications is a large buffer memory, due to the high rates of data generated. We investigate priority assignment algorithms for static-priority, preemptive real-time scheduling of periodic task sets. These algorithms are directed at applications in which worst case execution latency is not as important as buffer minimization. Examples of such applications include video and audio playout / recording, browsing through a database with audio / video data, and all types of non-interactive real-time applications. We refer to these as throughput-oriented real-time applications.

The techniques developed retain the simplicity of static priority scheduling, while improving the processor utilization which can be obtained. In addition, we are able to derive hard bounds on worst case execution time which are low enough for most practical deadline-based applications. We show that the standard rate-monotonic priority assignment algorithm is not optimal in terms of our goal of minimising input buffer size. Approximate algorithms are then found which perform better than the rate-monotonic algorithm, and bounds on their performance are derived. Average case performance is studied using simulation. The best algorithm combines rate-monotonic and shortest-job-first priority assignments. We show how to obtain deadline-based scheduling algorithms for task sets with arbitrary deadlines. The existence of an algorithm which minimizes buffer space is left as an open problem.

"Communication Networks for Autonomous Mobile Robot Computer Architectures and Distributed Real-Time Computing"

This thesis introduces an interprocessor communication network for mobile robot computer architectures and real-time communication networks. The objective of the research is to provide communication with very low delay.

The segmented bus is a reconfigurable interconnection network for intelligent mobile robot applications. The segmented bus employs a preemptive circuit-switched message transfer technique for low delay. For mobile robot applications, message delays are lower on the segmented bus than on alternatives such as multiple buses.

Interactive real-time applications require low end-to-end delay bounds and zero losses. This thesis offers solutions to admission control and low latency message transfers for interactive real-time message flows. Our proposed adaptive slack allocation method will accept a larger number of flows than a static allocation. Preemptive cut-through multiplexing reduces the minimum achievable end-to-end delay bounds to values close to that of circuit-switching. We show how these techniques can be applied to ATM networks, and assess the hardware and processing costs associated with these techniques.

"Low Latency, High Acceptance Real-Time Communication in Wide Area Networks"

This paper addresses the problem of providing very low guaranteed end-to-end delay bounds in packet-switched wide area networks such as B-ISDN. Such quality of service is necessary for real-time flows requiring no losses, and low end-to-end delay. These flows are used in continuous media, distributed sensing and command/control applications. Our work offers solutions to two problems: (i) flow admission control, and (ii) low-latency message transfers. The off-line flow admission control procedure adaptively reallocates slack to heavily loaded links. We show it results in higher acceptance rates than non-adaptive methods. The multiplexing or service method we propose combines preemption and cut-through to reduce the minimum achievable end-to-end delay bound. Preemptive cut-through can be implemented in ATM networks to provide tight delay bounds. The implementation cost of this method is examined as well.

"Routing and End-to-End Quality of Service in Multimedia Networks"

Packet / cell switched networks are increasingly being used to carry real-time multimedia traffic such as voice and video. This kind of traffic requires certain guarantees on the Quality of Service (QoS) it receives from the network. Typically, QoS measures include end-to-end latency, and loss probabilities. A key issue facing the network service provider is that of providing such QoS guarantees to users while utilizing network resources as efficiently as possible. In this thesis, three problems are analyzed on increasing the efficiency of network resource usage while providing strong guarantees on the end-to-end QoS received by users of such networks.

The first problem is whether existing routing algorithms can be used for real-time traffic. New routing algorithms are developed which perform better than algorithms not designed specifically for real-time traffic. It is shown that routing algorithms designed to exploit the constraints of the admission control algorithm result in improved call acceptance. The interaction between the routing and admission control functions is studied. The relative performance of three standard admission control schemes is also studied.

The second problem is what combination of multiplexing policies to use for efficient transport for real-time traffic. This thesis suggests an architecture which results in improved bandwidth utilization over existing schemes which provide end-to-end QoS guarantees. The architecture is based on deterministic bandwidth reservation at the level of Virtual Paths (i.e., on groups of channels rather than individual channels). It is shown that significant utilization levels are achievable while retaining the ability to provide end-to-end QoS guarantees. The problem of distributing given network resources over the different physical links traversed by a VP in the proposed architecture is addressed. The improved resource usage of an approach which lets all the loss occur at the first physical hop is proved. An implementation of the proposed architecture is suggested and examined.

The final problem addressed in this thesis is how to allocate resources with minimal user input and guidance. An algorithm is developed for determining on-line the minimum resources required to satisfy a specified QoS. The use of this technique is demonstrated for determining the minimum bandwidth allocation needed to meet a specified constraint on average loss probability. The algorithm is shown to be robust and accurate under different conditions while requiring very little knowledge of the traffic generation statistics of the source. The technique is shown to be applicable to different types of resources and different QoS definitions.

Dr. D. S. Reeves, reeves@eos.ncsu.edu

The Multicast Routing Simulator

- Tar File for the Multicast Routing Simulator, by Hussein Salama. Contains source code, make files, user manual, and sample data. Note that this source compiles on the Sun under SunOS, but not under Solaris. For machines running Solaris, you should upload the binary executable instead.(0.7 MB)
 - Binary executable for the Multicast Routing Simulator on the Decstation Architecture. (2.8 MB)
 - Binary executable for the Multicast Routing Simulator on the IBM RS6000 Architecture. (4.2 MB)
 - Binary executable for the Multicast Routing Simulator on the Sun Sparc Architecture. (2.0 MB)
-

Dr. D. S. Reeves, reeves@eos.ncsu.edu

The MPEG-2 Transport Stream Browser

- Tar File for the MPEG-2 Transport Stream Browser tool, by Mike Izquierdo. It allows users to navigate an MPEG-2 Transport Stream and inspect transport packet headers, adaptation fields and program association tables. Other commands such as find pid, find video pes, and find audio pes are included as well. This version runs on PCs using the Windows 3.1 OS.
 - Self-unpacking binary executable for the PC .
 - README file for m2tsbwin.
 - Tar file for a Unix Version of the MPEG-2 Transport Stream Browser.
 - README file for the Unix Version of the MPEG-2 Transport Stream Browser.
-

Dr. D. S. Reeves, reeves@eos.ncsu.edu